Mutative σ-Self-Adaptation Can Beat Cumulative Step Size Adaptation when Using Weighted Recombination

Hans-Georg Beyer
Research Center Process and Product Engineering
Department of Computer Science
Vorarlberg University of Applied Sciences
Achstr. 1, A-6850 Dornbirn, Austria
hans-georg.beyer@fhv.at

Alexander Melkozerov
Research Center Process and Product Engineering
Department of Computer Science
Vorarlberg University of Applied Sciences
Achstr. 1, A-6850 Dornbirn, Austria
alexander.melkozerov@fhv.at

ABSTRACT

This paper proposes the σ-self-adaptive weighted multirecombination evolution strategy (ES) and presents a performance analysis of this newly engineered ES. The steady state behavior of this strategy is investigated on the sphere model and a formula for the optimal choice of the learning parameter is derived allowing the ES to reach maximal performance. A comparison between weighted multirecombination ES with σ-self-adaptation (σSA) and with cumulative step size adaptation (CSA) shows that the σ-self-adaptive ES can exhibit the same performance and can even outperform its CSA counterpart for a range of learning parameters.

Categories and Subject Descriptors

I.2.6 [Artificial Intelligence]: Learning—parameter learning; G.1.6 [Numerical Analysis]: Optimization

General Terms

Algorithms, Design, Theory, Experimentation, Performance

Keywords

Weighted multirecombination, evolution strategy, mutation strength self-adaptation, cumulative step length adaptation

1. INTRODUCTION

Evolution strategies (ES) are a sub-class of nature-inspired direct search (and optimization) methods which use mutation, recombination, and selection applied to a population of individuals containing candidate solutions in order to evolve iteratively better and better solutions [6]. For each offspring individual, the canonical $(\mu/\rho \overset{+}{,} \lambda)$-ES performs mutation of strategy parameters and uses these to control the mutation of the offspring parameter vector. A control mechanism for such strategy parameters can provide an ability to adapt the ES's behavior to the features of particular objective function, thus, improving its overall performance aiming at a faster and more reliable approach to the optimizer state. As for ES in real-valued search spaces, the mutation strength σ is the most important strategy parameter that must be adapted continuously. It basically determines the ES's step length. Different adaptation techniques have been developed with the aim of obtaining optimal performance, including the 1/5-rule [14], self-adaptation [14, 15] and cumulative step length adaptation [13].

The mutation strength self-adaptation mechanism controls the mutation strength by means of evolution. Each ES individual contains, in addition to the parameter vector, its own mutation strength, which is mutated by multiplication with certain random numbers (often log-normally distributed). As the resulting mutation strength determines the standard deviation of the mutation vector, the algorithm performs the adaptation of the mutation strength indirectly due to selection. Selection itself is based on the objective function values calculated from the offspring parameter vectors. Since no external mechanisms are required for the adaptation of the mutation strength, this version of the adaptation is referred to as σ-self-adaptation (σSA). A quantitative analysis of σSA can be found in [12].

As an alternative, the cumulative step length adaptation (CSA) has been proposed [13]. The CSA is based on the assumption that uncorrelated consecutive search steps of the ES correspond to optimally chosen mutation strength. CSA maintains a fading memory of search path and continously compares its length with the expected length of an "ideal" path consisting of random steps. The mutation strength is changed in accordance with the result of this comparison. A discussion considering the optimality of the basic assumption in CSA was presented in [7].

Recently, a modified CSA version has been proposed by Arnold [1] as a mutation strength adaptation procedure for weighted multirecombination evolution strategies which perform a weighted multirecombination of all λ offspring individuals in a population. As has been shown for the sphere model, CSA allows for a fixed choice of weights, which guarantees a nearly optimal performance (up to a factor of $\sqrt{2}-1$ from a theoretical prediction) on the sphere model.

While the standard CSA can drive the ES to nearly optimal performance (maximal progress rate) on the sphere

GECCO'08, July 12–16, 2008, Atlanta, Georgia, USA.

model in a robust fashion, the alternative standard σSA exhibits a sensitive dependency of performance on the learning parameter [8].

In this paper, it will be shown that one can design a self-adaptive ES that performs comparably well when using weighted multirecombination on the object parameters. We will analyze this new ES and compare its performance with the weighted multirecombination ES with CSA.

The remaining part of this paper is organized as follows. Section 2 describes the $(\mu/\mu_I, \lambda)$-σSA-ES, the weighted multirecombination ES with CSA and a newly engineered σ-self-adaptive weighted multirecombination ES. Section 3 introduces ES progress measures and steady state formulas which lead to an optimal self-adaptation rate formula. In Section 4, comparison with experiments in finite-dimensional search spaces is presented. Section 5 contains a brief summary of the results obtained and outlines future research steps.

2. ENGINEERING THE NEW $(\lambda)_{\mathrm{OPT}}$-σ-SELF-ADAPTATION ES

This section describes the $(\mu/\mu_I, \lambda)$-σSA-ES, the weighted multirecombination ES with CSA and the σ-self-adaptive weighted multirecombination ES, each of which is working with populations of λ offspring individuals. Each individual can be defined as [6]

$$\text{ES individual} \quad \mathbf{a} := (\mathbf{y}, \mathbf{s}, f(\mathbf{y})), \tag{1}$$

where $\mathbf{y}$ is the parameter vector, $\mathbf{s}$ is a set of strategy parameters and $f(\mathbf{y})$ is the individual's value of the objective function to be optimized.

2.1 The $(\mu/\mu_I, \lambda)$-σ-self-adaptation-ES

This ES uses all μ parental individuals for creation of an offspring. The subscript I denotes the intermediate recombination, which calculates a recombinant individual as the centroid of all μ parental individuals, e.g., as average values of the parental parameters vectors and strategy parameters, respectively.

The $(\mu/\mu_I, \lambda)$-σ-self-adaptation-ES is given below:

1. Initialize parent state

$$\begin{aligned} \sigma_{\mathrm{p}} &\leftarrow \sigma_{\mathrm{init}} \\ \mathbf{y}_{\mathrm{p}} &\leftarrow \mathbf{y}_{\mathrm{init}} \end{aligned}$$

2. Generate λ offspring according to

$$\forall l = 1, \ldots, \lambda: \quad \begin{cases} \tilde{\sigma}_l \leftarrow \sigma_{\mathrm{p}} \mathrm{e}^{\tau \mathcal{N}_l(0,1)}, \\ \tilde{\mathbf{z}}_l \leftarrow \boldsymbol{\mathcal{N}}_l(\mathbf{0}, \mathbf{I}), \\ \tilde{\mathbf{y}}_l \leftarrow \mathbf{y}_{\mathrm{p}} + \tilde{\sigma}_l \tilde{\mathbf{z}}_l, \\ \tilde{\mathrm{f}}_l \leftarrow f(\tilde{\mathbf{y}}_l). \end{cases}$$

   where $\mathcal{N}_l(0,1)$ is a $(0,1)$ normally distributed random scalar, $\boldsymbol{\mathcal{N}}_l(\mathbf{0}, \mathbf{I})$ is a $(0,1)$ normally distributed random vector, $\tilde{\sigma}_l$ is the mutation strength, and $\tilde{\sigma}_l \tilde{\mathbf{z}}_l$ is the so-called mutation vector.

3. Order the λ offspring according to its objective function values.

4. Perform recombination of mutation strengths and parameter vectors according to

$$\langle \sigma \rangle \leftarrow \frac{1}{\mu} \sum_{m=1}^{\mu} \tilde{\sigma}_{m;\lambda} \tag{2}$$

$$\langle \mathbf{y} \rangle \leftarrow \frac{1}{\mu} \sum_{m=1}^{\mu} \tilde{\mathbf{y}}_{m;\lambda}, \tag{3}$$

   where subscript "$m;\lambda$" denotes the mth best offspring individual generated during step 2, e.g., the individual with the mth-smallest value of the objective function $f(\mathbf{y})$ for minimization.

5. Create new parent state

$$\begin{aligned} \sigma_{\mathrm{p}} &\leftarrow \langle \sigma \rangle \\ \mathbf{y}_{\mathrm{p}} &\leftarrow \langle \mathbf{y} \rangle \end{aligned} \tag{4}$$

6. Goto 2. until termination criterion fulfilled.

The change of the mutation strength is implemented in the algorithm using the so-called log-normal operator $\mathrm{e}^{\tau \mathcal{N}_l(0,1)}$. The learning parameter τ in the log-normal operator controls the self-adaptation rate. Theoretical and empirical investigations have shown that in order to get optimal linear stochastic convergence on the sphere model τ must be chosen proportional to $1/\sqrt{N}$, i.e.

$$\tau = \frac{\alpha}{\sqrt{N}}. \tag{5}$$

As for the sphere model without noise, $N \to \infty$, and large populations, a good choice of α is given by [11]

$$\alpha = \frac{1}{\sqrt{2}}. \tag{6}$$

After the termination criterion is fulfilled, the current parent state is considered as an approximation of the optimizer of the objective function $f(\mathbf{y})$.

2.2 The $(\lambda)_{\mathrm{opt}}$-ES

A known disadvantage of all $(\mu/\rho \overset{+}{,} \lambda)$-ES is that they are not using all information about the offspring individuals generated during each ES generation [1]. In order to overcome this weak point, a weighted multirecombination can be used [1] which takes into account information about the ranking of all offspring individuals. That is, it does not discard the worst $(\lambda-\mu)$ individuals like the $(\mu/\rho \overset{+}{,} \lambda)$-ES does. The influence of each individual is determined by its ranking order which is based on the objective function values.

The weighted multirecombination ES can be described in the following way. It creates λ new offspring according to

$$\forall l = 1, \ldots, \lambda: \quad \mathbf{y}_l \leftarrow \mathbf{y}_{\mathrm{p}} + \sigma_l \boldsymbol{\mathcal{N}}_l(\mathbf{0}, \mathbf{I}), \tag{7}$$

where σ_l is determined by a mutation strength adaptation mechanism. In CSA, $\sigma_l = \sigma$, i.e., this adaptation mechanism needs only one mutation strength per generation.

After offspring procreation, the weighted multirecombination ES calculates the corresponding objective function values $f(\mathbf{y})$. Thereafter, the ES ranks the individuals w.r.t. their fitnesses (objective function values) and computes the weighted sum

$$\langle \mathbf{z} \rangle_\omega \leftarrow \sum_{l=1}^{\lambda} \omega_{l,\lambda} \mathbf{z}^{(l;\lambda)} \tag{8}$$

of the vectors $\mathbf{z}$. The superscript $(l;\lambda)$ refers to the lth-best of the λ offspring (the lth-smallest for minimization). Weights $\omega_{l,\lambda}$ are dependent on the rank of the individual in the set of all offspring individuals [1].

A new parent state is obtained as

$$\mathbf{y}_{\mathrm{p}} \leftarrow \mathbf{y}_{\mathrm{p}} + \langle\sigma\rangle\langle\mathbf{z}\rangle_\omega, \qquad (9)$$

where $\langle\sigma\rangle = \sigma$ for the CSA adaptation mechanism.[1]

In order to derive a condition for optimal choice of weights $\omega_{l,\lambda}$ for the sphere model, it is necessary to define a performance measure for the weighted multirecombination ES. One option is to use the expected change of the objective function value $f(\mathbf{y})$ from one generation to the next, which is referred to as quality gain Δ. Thus, optimality of weights $\omega_{l,\lambda}$ can be defined as those weights which maximize the quality gain of the ES.

In this work, we will refer to the weighted multirecombination evolution strategy with optimally chosen weights $\omega_{l,\lambda}$ as $(\lambda)_{\mathrm{opt}}$-ES.

2.3 Building the $(\lambda)_{\mathrm{opt}}$-σ-Self-Adaptation-ES

In this section we design the new weighted ES with σSA. It is the result of the combination of the σSA mechanism with the $(\lambda)_{\mathrm{opt}}$-ES. In order to use in this case theoretical results derived in [12] for the $(\mu/\mu_I,\lambda)$-σSA-ES, we construct the algorithm of the $(\lambda)_{\mathrm{opt}}$-σSA-ES in such a manner, that it incorporates the weighted multirecombination with minimal changes of the original $(\mu/\mu_I,\lambda)$-σSA-ES algorithm.

The algorithm of the $(\lambda)_{\mathrm{opt}}$-σSA-ES reads:

1. Initialize the parent state

$$\begin{aligned}\sigma_{\mathrm{p}} &\leftarrow \sigma_{\mathrm{init}}\\ \mathbf{y}_{\mathrm{p}} &\leftarrow \mathbf{y}_{init}\end{aligned}$$

2. Generate λ offspring according to

$$\forall l = 1,\dots,\lambda: \begin{cases}\tilde{\sigma}_l \leftarrow \sigma_{\mathrm{p}} \mathrm{e}^{\tau\mathcal{N}_l(0,1)},\\ \tilde{\mathbf{z}}_l \leftarrow \boldsymbol{\mathcal{N}}_l(\mathbf{0},\mathbf{I}),\\ \tilde{\mathbf{y}}_l \leftarrow \mathbf{y}_{\mathrm{p}} + \tilde{\sigma}_l\tilde{\mathbf{z}}_l,\\ \tilde{f}_l \leftarrow f(\tilde{\mathbf{y}}_l).\end{cases}$$

3. Rank λ offspring according to their f-values.

4. Perform recombination of mutation strengths

$$\langle\sigma\rangle \leftarrow \frac{1}{\mu}\sum_{m=1}^{\mu}\tilde{\sigma}_{m;\lambda}. \qquad (10)$$

5. Compute the weighted sum $\langle\mathbf{z}\rangle_\omega$ of mutation vectors

$$\langle\mathbf{z}\rangle_\omega \leftarrow \sum_{l=1}^{\lambda}\omega_{l;\lambda}\tilde{\mathbf{z}}_{l;\lambda}. \qquad (11)$$

6. Create new parent state

$$\begin{aligned}\sigma_{\mathrm{p}} &\leftarrow \langle\sigma\rangle\\ \mathbf{y}_{\mathrm{p}} &\leftarrow \mathbf{y}_{\mathrm{p}} + \langle\sigma\rangle\langle\mathbf{z}\rangle_\omega\end{aligned}$$

7. Goto 2. until termination criterion fulfilled.

[1]Note, in Eq. (9) already a generalization has been introduced in that we used $\langle\sigma\rangle$ instead of σ which is used in CSA.

Comparing this algorithm with the $(\mu/\mu_I,\lambda)$-σSA-ES, one notes a difference in the way how parameter recombination is performed in Step 5 involving the calculation of the weighted sum $\langle\mathbf{z}\rangle_\omega$. The weighted sum $\langle\mathbf{z}\rangle_\omega$ is used to determine the direction to the new parental state. This is in contrast to the centroid value $\langle\mathbf{y}\rangle$ calculation in the conventional $(\mu/\mu_I,\lambda)$-ES algorithm. At the same time, no changes of the σSA mechanism are introduced. This circumstance allows us to apply theoretical formulas, derived for the $(\mu/\mu_I,\lambda)$-σSA-ES in [12], for the analysis of the $(\lambda)_{\mathrm{opt}}$-σSA-ES. The analysis will be presented in Section 3.1. We will use μ as the number of individuals participating in the calculation of the mutation strength recombinant $\langle\sigma\rangle$ when describing the actual $(\lambda)_{\mathrm{opt}}$-σSA-ES.

3. PERFORMANCE ANALYSIS OF THE $(\lambda)_{\mathrm{OPT}}$-σSA-ES

In this work, the analysis of the ES is performed on the sphere model. Without loss of generality, the quadratic sphere is considered. It is one of the commonly used test functions for unconstrained optimization [4]. This function maps candidate solution $\mathbf{y}$ to the square of its Euclidean distance $r = \|\hat{\mathbf{y}} - \mathbf{y}\|$ from the optimizer $\hat{\mathbf{y}} \in \mathbb{R}^N$, where N is the search space dimensionality. It reads

$$f(\mathbf{y}) = \|\hat{\mathbf{y}} - \mathbf{y}\|^2,$$

where the task is minimization. The sphere serves as a model for objective functions in the vicinity of well-behaved local optima. Using the sphere environment, we can derive important theoretical predictions, which – sometimes – can be extended to other types of objective functions. Furthermore, ES with covariance matrix adaptation (CMA-ES) have been found to effectively transform a wide range of convex quadratic functions into the sphere [9].

In the next two sections, we will shortly review the results of theoretical analyses obtained for the $(\mu/\mu_I,\lambda)$-σSA-ES and the $(\lambda)_{\mathrm{opt}}$-ES. Furthermore, we provide preparatory steps which will finally be used in Section 3.3 to derive our main theoretical result on the optimal choice of the learning parameter.

3.1 Analysis of $(\mu/\mu_I,\lambda)$-σSA-ES

Let $r^{(g)}$ denote the distance of the centroid in the gth generation $\langle\mathbf{y}\rangle^{(g)}$ to the optimum $\hat{\mathbf{y}}$, i.e., $r^{(g)} = \|\mathbf{r}^{(g)}\| = \|\langle\mathbf{y}\rangle^{(g)} - \hat{\mathbf{y}}\|$ and $s^{(g)} = \langle\sigma^{(g)}\rangle$ the mean value of the parental σ-values. The central quantities describing the behavior of the ES are the progress rate φ and the self-adaptation response.

The expected change of the distance r from one generation to the next defines the progress rate

$$\varphi(s^{(g)}, r^{(g)}) = \mathrm{E}\left[r^{(g)} - r^{(g+1)} \;\middle|\; s^{(g)}, r^{(g)}\right]. \qquad (12)$$

Considering the mutation strength, the expected relative change is called the self-adaptation response (SAR)

$$\psi(s^{(g)}, r^{(g)}) = \mathrm{E}\left[\frac{s^{(g+1)} - s^{(g)}}{s^{(g)}} \;\middle|\; s^{(g)}, r^{(g)}\right]. \qquad (13)$$

SAR provides information on the self-adaptation mechanism feedback resulting in the change of the mutation strength in generation $g+1$. This feedback is determined solely by the internal state of the ES at generation g. The learning parameter τ controls the magnitude of the SAR feedback: increasing τ allows to obtain larger mutation strength changes for a given ES state since the exponent parameter $\tau\mathcal{N}_l(0,1)$ in the log-normal operator grows and newly created offspring get mutation strengths which differ from $s^{(g)}$ by larger values. Ideally, the SAR function should be positive for mutation strength smaller than the optimal one and negative otherwise driving in that way the mutation strength to an optimal value. In order to obtain such an ideal behavior, it is necessary to choose the learning parameter τ appropriately taking the search space dimensionality into account. This also requires a certain knowledge about the fitness landscape which is usually unknown in practice. Approximate formulas for optimal τ depending on the ES' exogenous strategy parameters can be used in this case.

In order to obtain state variables that are independent of the position in the search space, one uses the normalizations

$$\varphi^* = \varphi\frac{N}{r^{(g)}} \tag{14}$$

and

$$s^{*(g)} = s^{(g)}\frac{N}{r^{(g)}}. \tag{15}$$

Using several simplifications (consideration of the limit case $\tau \to 0$, the asymptotic behavior for $N \to \infty$, and Taylor series expansions), an approximate formula for the self-adaptation response can be derived [12]

$$\psi(s^{*(g)}) \approx \tau^2\left(\frac{1}{2} + e^{1,1}_{\mu,\lambda} - s^{*(g)}c_{\mu/\mu,\lambda}\right). \tag{16}$$

The so-called generalized progress coefficients $e^{a,b}_{\mu,\lambda}$ used in (16) are given by

$$e^{a,b}_{\mu,\lambda} = \frac{\lambda-\mu}{\sqrt{2\pi}^{a+1}}\binom{\lambda}{\mu}\int_{-\infty}^{\infty} t^b \mathrm{e}^{-\frac{a+1}{2}t^2}\Phi(t)^{\lambda-\mu-1}(1-\Phi(t))^{\mu-a}\mathrm{d}t, \tag{17}$$

where $\Phi(t)$ is the cumulative distribution function of the standard normal variate. The progress coefficient $c_{\mu/\mu,\lambda}$ is a special case of the generalized progress coefficients

$$c_{\mu/\mu,\lambda} = e^{1,0}_{\mu,\lambda}.$$

3.2 Analysis of the $(\lambda)_{\mathrm{opt}}$-ES

While we need the SAR results from the $(\mu/\mu_I,\lambda)$-ES, the progress of the new $(\lambda)_{\mathrm{opt}}$-σSA-ES in the object parameter space must be taken from the analysis of the $(\lambda)_{\mathrm{opt}}$-ES. The expected difference of two consecutive parental fitness values is referred to as quality gain. It is defined as

$$\Delta = \mathrm{E}\left[f(\langle\mathbf{y}\rangle^{(g)}) - f(\mathbf{y}^{(g+1)})\right]. \tag{18}$$

We now make plausible that the normalized quality gain

$$\Delta^* = \Delta\frac{N}{2\left(r^{(g)}\right)^2} \tag{19}$$

formally agrees with the formula of the normalized progress rate (14) for the quadratic sphere environment in the limit $N \to \infty$ [3].

For finite normalized mutation strength, the normalized fitness gain is finite. Therefore, with increasing N, the expected value of the difference between the distance $r^{(g)}$ of the parental centroid from the location of the optimum and the distance $r^{(g+1)}$ of the selected offspring to the location of the optimum goes to zero. Since

$$\begin{aligned}\Delta^* &\overset{N\to\infty}{=} \frac{N}{2\left(r^{(g)}\right)^2}\mathrm{E}\left[f(\langle\mathbf{y}\rangle^{(g)}) - f(\mathbf{y}^{(g+1)})\right]\\ &= \mathrm{E}\left[\frac{N}{2\left(r^{(g)}\right)^2}\left((r^{(g)})^2 - (r^{(g+1)})^2\right)\right]\\ &= \mathrm{E}\left[N\frac{r^{(g)}-r^{(g+1)}}{r^{(g)}}\cdot\frac{r^{(g)}+r^{(g+1)}}{2r^{(g)}}\right],\end{aligned} \tag{20}$$

and as $\frac{r^{(g)}+r^{(g+1)}}{2r^{(g)}}$ tends to one, it follows

$$\Delta^*(s^{*(g)}) \overset{N\to\infty}{=} \mathrm{E}\left[\frac{N}{r^{(g)}}\left(r^{(g)} - r^{(g+1)}\right)\cdot 1\right] = \varphi^*(s^{*(g)}). \tag{21}$$

This equation will be needed below to derive the steady state behavior.

In order to use (21), the quality gain Δ^* using optimally chosen weights $\omega_{k;\lambda}$ is needed. Using several simplifications (consideration of the asymptotic behavior for $N \to \infty$, assumption that the normalized mutation strength σ^* is of $\mathcal{O}(1)$, Taylor series expansion), an asymptotically exact formula for the normalized quality gain of the $(\lambda)_{\mathrm{opt}}$-ES for optimal weights

$$\omega_{k,\lambda} = E_{k,\lambda} \text{ for } k = 1,\ldots,\lambda, \tag{22}$$

has been derived in [1]

$$\Delta^*(\sigma^{*(g)}) = W_\lambda\left(\sigma^{*(g)} - \frac{(\sigma^{*(g)})^2}{2}\right). \tag{23}$$

The W_λ is defined as

$$W_\lambda = \sum_{k=1}^{\lambda} E^2_{k,\lambda}, \tag{24}$$

where $E_{k,\lambda}$ denotes the expectation of the $(\lambda+1-k)$th order statistic of the standard normal variate. This is another special case of the generalized progress coefficients (17)

$$E_{k,\lambda} = e^{0,1}_{k-1,\lambda}. \tag{25}$$

It is important to note that Eq. (23) was derived considering a single normalized mutation strength $\sigma^{(g)}$ for each generation g. From the algorithmic point of view, this does not hold for the $(\lambda)_{\mathrm{opt}}$-σSA-ES since each offspring individual has its own mutation strength $\tilde{\sigma}_l$. However, from the theoretical point of view, considering the asymptotical limit case $N \to \infty$, we will show that one can assume that the mean value $s^{(g)}$ can be used in Eq. (23) instead of $\sigma^{(g)}$. Since

$$s^{(g)} = \langle\sigma^{(g)}\rangle = \frac{1}{\mu}\sum_{m=1}^{\mu}\sigma^{(g)}_{m;\lambda} = \frac{1}{\mu}\sum_{m=1}^{\mu}\sigma^{(g-1)}_{\mathrm{p}}\mathrm{e}^{\tau\mathcal{N}_{m;\lambda}(0,1)}$$

and taking (5) into account

$$\tau = \frac{\alpha}{\sqrt{N}} \overset{N\to\infty}{\longrightarrow} 0,$$

it follows that

$$s^{(g)} \overset{N\to\infty}{=} \frac{1}{\mu}\sum_{m=1}^{\mu}\sigma^{(g-1)}_{\mathrm{p}}\cdot 1 = \sigma^{(g-1)}_{\mathrm{p}}. \tag{26}$$

The mean value $s^{(g)}$ of mutation strengths generated using the log-normal operator in Step 2 of the $(\lambda)_{\text{opt}}$-σSA-ES is asymptotically equal to the parental mutation strength $\sigma_{\text{p}}^{(g-1)}$ in the limit $N \to \infty$. This transfers also to the normalized quantities. Thus, we can use $s^{*(g)}$ in Eq. (23) instead of $\sigma^{*(g)}$.

3.3 Analyzing the Steady State of the $(\lambda)_{\text{opt}}$-σSA-ES

In the case of a correctly working σSA, the expected value of the normalized mutation strength reaches a stationary state over time with $\lim_{g\to\infty} s^{*(g)} = s^*_{\text{st}}$. Using the deterministic equilibrium condition [12]

$$s^{*(g+1)} = s^{*(g)} \frac{1 + \psi(s^{*(g)}, N)}{1 - \frac{\varphi^*(s^{*(g)}, N)}{N}}, \tag{27}$$

one obtains the steady state condition

$$\frac{\varphi^*(s^{*(g)}_{\text{st}})}{N} = -\psi\left(s^{*(g)}_{\text{st}}\right) \tag{28}$$

which will be used to derive the stationary mutation strength. Due to the asymptotic equality (21) $\varphi^*(s^{*(g)}) = \Delta^*(s^{*(g)})$, which is correct for the quadratic sphere environment in the limit of the infinite parameter space dimensionality $N \to \infty$, one can use (23) instead of $\varphi^*(s^{*(g)})$ in (28). Inserting (16) and (23) into (28) leads to

$$W_\lambda \left(s^*_{\text{st}} - \frac{(s^*_{\text{st}})^2}{2}\right) = -\alpha^2 \left[\frac{1}{2} + e^{1,1}_{\mu,\lambda} - s^*_{\text{st}} c_{\mu/\mu,\lambda}\right], \tag{29}$$

where the learning parameter τ in Eq. (16) has been chosen according to (5). Thus, Eq. (16) becomes independent of the search space dimensionality N.

As one can see, the steady state depends on α. Provided that α is given, s^*_{st} can be calculated analytically. Solving the quadratic equation (29) for s^*_{st} yields

$$s^*_{\text{st}} = 1 - \frac{c_{\mu/\mu,\lambda}\alpha^2}{W_\lambda} + K, \tag{30}$$

where

$$K = \sqrt{1 + \left(1 - 2c_{\mu/\mu,\lambda} + 2e^{1,1}_{\mu,\lambda}\right)\frac{\alpha^2}{W_\lambda} + \frac{c^2_{\mu/\mu,\lambda}\alpha^4}{W^2_\lambda}}. \tag{31}$$

Let us discuss the extreme choices of the N-independent factor α used in the learning parameter formula (5). As we will see, the resulting steady state mutation strengths correspond to the characteristic zeros of the SAR function and the quality gain, respectively. Dividing (29) by $-\alpha^2$ and taking the limit $\alpha \to \infty$ (provided that $s^*_{\text{st}} < \infty$), one gets

$$\lim_{\alpha\to\infty} -\frac{1}{\alpha^2} W_\lambda \left(s^*_{\text{st}} - \frac{(s^*_{\text{st}})^2}{2}\right) = 0. \tag{32}$$

Therefore, the bracket on the right hand side of (29) must vanish in the limit $\alpha \to \infty$. Thus, the steady state mutation strength s^*_{st} is determined by the zero of the SAR, i.e., $s^*_{\text{st}} = s^*_{\psi_0}$. In order to obtain $s^*_{\psi_0}$, one has to solve

$$\frac{1}{2} + e^{1,1}_{\mu,\lambda} - s^*_{\psi_0} c_{\mu/\mu,\lambda} = 0$$

for $s^*_{\psi_0}$ resulting in

$$s^*_{\psi_0} = \frac{\frac{1}{2} + e^{1,1}_{\mu,\lambda}}{c_{\mu/\mu,\lambda}}. \tag{33}$$

The other extreme case is given by $\alpha \to 0$. Using Eq. (31) and (30), the limit value becomes

$$\lim_{\alpha\to 0} s^*_{\text{st}} = 2,$$

Comparing with (29), one sees that this is the second zero of the quality gain (23)

$$s^*_{\Delta_0} = 2. \tag{34}$$

Considering the two extreme choices of α resulting in the characteristic zeros $s^*_{\psi_0}$ (33) and $s^*_{\Delta_0} = 2$ (34), one can conjecture that for non-extreme choices of α

$$s^*_{\text{st}} \in (s^*_{\psi_0}, s^*_{\Delta_0}) \tag{35}$$

(provided that $s^*_{\psi_0} < s^*_{\Delta_0}$). This is indeed the case and can easily be shown then looking at Fig. 3: The steady state s_{st} is determined by the intersection of the (linear) SAR function (multiplied by $-N$) and the quadratic φ^* function. The angle of inclination of the SAR function depends on α. If $\alpha = 0$ then slope of SAR is zero resulting in $s^*_{\text{st}} = s^*_{\Delta_0}$. Increasing α gradually, the point of intersection shifts to smaller s^*_{st}-values, i.e., $s^*_{\text{st}}(\alpha)$ is a monotonously decreasing function of α. Finally, the SAR ends up as a vertical line yielding the zero of the SAR $s^*_{\psi_0}$.

Now, let us consider the steady state progress rate as a function of α. If we take into account the equality $\varphi^*(s^{*(g)}) = \Delta^*(s^{*(g)})$ and insert the stationary mutation strength (30) into the quality gain (23), we obtain the stationary progress rate

$$\varphi^*_{\text{st}}(\alpha) = \frac{W_\lambda}{2}\left[1 - \left(\frac{c_{\mu/\mu,\lambda}\alpha^2}{W_\lambda} - K\right)^2\right]. \tag{36}$$

Considering (23), one easily sees that the quality gain reaches its maximal value $\Delta^*_{\max} = W_\lambda/2$ at $s^*_{\Delta_{\max}} = 1$. Provided that $s^*_{\Delta_{\max}}$ is in the interval (35), one can tune the ES for maximal performance using α. As for maximal progress in the stationary state, one would like to have $s^*_{\text{st}} = s^*_{\Delta_{\max}} = 1$. Using (30), this leads to

$$1 - \frac{c_{\mu/\mu,\lambda}\alpha^2_{\text{opt}}}{W_\lambda} + \sqrt{1 + \left(1 - 2c_{\mu/\mu,\lambda} + 2e^{1,1}_{\mu,\lambda}\right)\frac{\alpha^2_{\text{opt}}}{W_\lambda} + \frac{c^2_{\mu/\mu,\lambda}\alpha^4_{\text{opt}}}{W^2_\lambda}} = 1.$$

Solving for α_{opt}, one finally obtains

$$s^*_{\psi_0} < 1: \quad \alpha_{\text{opt}} = \sqrt{\frac{W_\lambda}{2c_{\mu/\mu,\lambda} - 2e^{1,1}_{\mu,\lambda} - 1}}. \tag{37}$$

Note, if $s^*_{\psi_0} > 1$ (cf. Eq. 33), then the denominator in (37) gets negative and the ES cannot be tuned to the theoretical quality gain maximum $W_\lambda/2$. This is the reason why truncation ratios $\mu/\lambda \geq 0.3$ should be chosen. Truncation ratios $\mu/\lambda = [0.3, 0.4]$ may be regarded as reasonable choices.

A list of α_{opt} values using Eq. (37) can be found in Table 1. As a rule of thumb, using large $\lambda \geq 50$, the following approximations may be used for α_{opt}

$$\mu/\lambda = 0.3: \quad \alpha_{\text{opt}} \approx 3\sqrt{\lambda},$$
$$\mu/\lambda = 0.4: \quad \alpha_{\text{opt}} \approx \frac{3}{2}\sqrt{\lambda}.$$

Table 1: Optimal parameter α values

(μ,λ)	$(3,10)$	$(15,50)$	$(30,100)$	$(300,1000)$
α_{opt}	8.6	21	31	99
(μ,λ)	$(4,10)$	$(20,50)$	$(40,100)$	$(400,1000)$
α_{opt}	4.6	11	15	48

4. COMPARISON WITH EXPERIMENTS

The derivation of the stationary state formulas was carried out in the limit of infinite search space dimensionality. These formulas have been derived using the simplifications described in Section 3. For finite N, it is necessary to verify the derivations presented above. Numerical experiments are used in this section in order to verify the derived formulas in finite-dimensional search spaces.

4.1 Performance of the $(\lambda)_{\mathrm{opt}}$-σSA-ES in Finite-Dimensional Search Spaces

The predictions of (36) are compared with the results of experiments in Fig. 1. A start vector $\mathbf{y}^{(0)} = \mathbf{1000}$ and an initial mutation strength $\sigma^{(0)} = 1$ were used for the $(\lambda)_{\mathrm{opt}}$-σSA-ES with $\lambda = 10$ offspring. The strategy parameter recombination was performed with $\mu = 4$ parents.

The sampling process was started from generation g_0 and executed until generation g. In order to calculate the stationary progress rate, the following formula [5] was used:

$$\varphi^*_{\mathrm{st}} = \frac{N}{g-g_0}\ln\left(\frac{r^{(g_0)}}{r^{(g)}}\right). \tag{38}$$

The simulations show that the normalized stationary progress

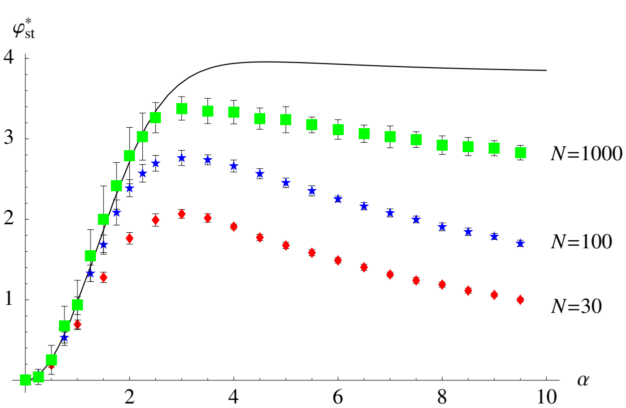

Figure 1: The stationary progress rate of the $(\lambda)_{\mathrm{opt}}$-σSA-ES ($\mu = 4$, $\lambda = 10$) as a function of the parameter α. The solid curve represent the result of (36). The points indicate the results of experiments averaged over 30 runs. Shown are from bottom to top the results of experiments for $N = 30$, $N = 100$, and $N = 1000$.

rate approaches the theoretical curve ($N \to \infty$) for smaller values of the parameter α. For larger α one can observe a degradation of the experimentally obtained normalized stationary progress rate due to the increase of the learning parameter τ, Eq. (5). Using a large learning parameter has advantages and disadvantages: A Larger learning parameter can reduce the transient time for the self-adaptation process, since it gives rise to larger SAR values for a given s^*. On the other hand, an increase in the learning parameter does also decrease the maximal attainable progress rate [5]. That is why, α should be chosen a certain percentage smaller than α_{opt} obtained for the infinite limit case $N \to \infty$.

4.2 Comparison with $(\mu/\mu_I,\lambda)$-σ-ES, $(\lambda)_{\mathrm{opt}}$-CSA-ES and Discussion

In order to compare the different strategies including the CSA-ES, we first recall results found in literature. The stationary progress rate formula for the self-adaptive $(\mu/\mu_I,\lambda)$-ES can be found in [12]. It reads

$$\varphi^*_{\mathrm{st}}(\alpha) = \alpha^2\left(\mu c^2_{\mu/\mu,\lambda}(1-\alpha^2) + c_{\mu/\mu,\lambda}K - \frac{1}{2} - e^{1,1}_{\mu,\lambda}\right), \tag{39}$$

where

$$K = \sqrt{\mu^2 c^2_{\mu/\mu,\lambda}\left(1-\alpha^2\right)^2 + 2\mu\alpha^2\left(\frac{1}{2} + e^{1,1}_{\mu,\lambda}\right)}.$$

As to the CSA-ES, we first present the cumulation rule. According to [1], the CSA is implemented in the $(\lambda)_{\mathrm{opt}}$-ES in the following way. Weighted sums $\langle\mathbf{z}\rangle_\omega$ are cumulated during the run of the ES in order to track steps in the search space using an N-dimensional vector $\mathbf{l}$. This vector is initialized according to $\mathbf{l}^{(0)} = \mathbf{0}$ and updated according to

$$\mathbf{l}^{(g+1)} = (1-c)\mathbf{l}^{(g)} + \sqrt{\frac{c(2-c)}{W_\lambda}}\langle\mathbf{z}\rangle^{(g)}_\omega, \tag{40}$$

where c is the cumulation parameter, $c = 1/\sqrt{N}$ [1]. The mutation strength is adapted using Arnold's update rule

$$\sigma^{(g+1)} = \sigma^{(g)}\mathrm{e}^{\frac{\|\mathbf{l}^{(g+1)}\|^2 - N}{2DN}}, \tag{41}$$

where D is the damping parameter, $D = 1/c$ [1]. Given these update rules, the steady state average quality gain of the $(\lambda)_{\mathrm{opt}}$-ES can be determined [1]

$$\Delta^*_{avg} = \left(\sqrt{2}-1\right)W_\lambda. \tag{42}$$

Due to the asymptotic equality $\varphi^*(s^{*(g)}) = \Delta^*(s^{*(g)})$, one can use the formula (42) for the comparison with the experimental results concerning the weighted multirecombination evolution strategies with σSA.

The outcome of (36), (39), and (42) are compared in Fig. 2 with results of experiments for the weighted multirecombination ES with σSA and CSA. Also depicted are the results of experiments for the weighted multirecombination ES with σSA using the particular choice of weights

$$\omega_{k,\lambda} = \begin{cases} 1/\mu & \text{if } 1 \leqq k \leqq \mu, \\ 0 & \text{otherwise} \end{cases} \tag{43}$$

In this case, the weighted multirecombination ES is simply the $(\mu/\mu_I,\lambda)$-σSA-ES. That is, the parental state $\mathbf{y}_{\mathrm{p}}$ is the centroid of the population that consists of the μ best of the λ offspring candidate solutions generated (cf Section 2.1).

The comparison of the $(\lambda)_{\mathrm{opt}}$-σSA-ES with the $(\lambda)_{\mathrm{opt}}$-CSA-ES shows that the weighted multirecombination ES with σSA can exhibit the same performance as the weighted multirecombination ES with CSA if the value of parameter

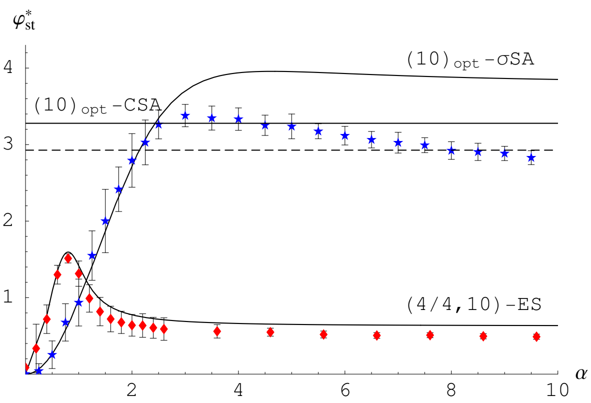

Figure 2: The stationary progress rate as a function of α. The solid horizontal line represents the theoretical result of (42) for the $(10)_{\text{opt}}$-CSA-ES, whereas the dashed depicts the results of respective experiments of the same $(10)_{\text{opt}}$-CSA-ES averaged over 30 runs (the standard deviation of the results of experiments is less than 0.05). The solid curves represent the results of (39) for the self-adaptive $(4/4_I, 10)$-ES and (36) for the $(\lambda)_{\text{opt}}$-σSA-ES. The points indicate the result of experiments for $N = 1000$ for the $(\lambda)_{\text{opt}}$-σSA-ES averaged over 30 runs. Shown are the results for the particular choice of weights (43), denoted by $(4/4, 10)$-ES, and for the optimal weights (22) $\omega_{k,\lambda} = E_{k,\lambda}$, denoted by $(10)_{\text{opt}}$-σSA.

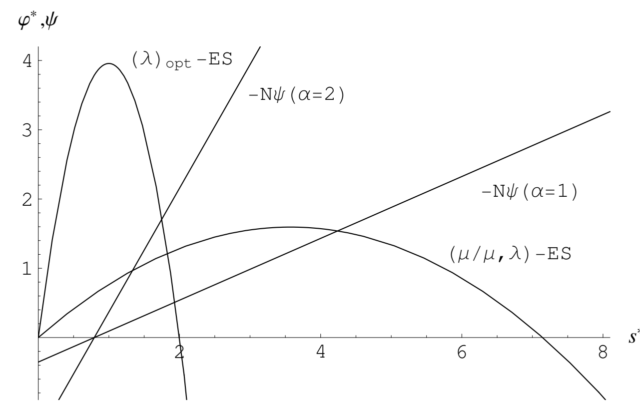

Figure 3: Graphical interpretation of the steady state conditions (28) for $(\mu/\mu_I, \lambda)$-ES and $(\lambda)_{\text{opt}}$-σSA-ES. Straight lines represent the SAR function ψ (16) with $\tau = \alpha/\sqrt{N}$ multiplied with $-N$ for $\alpha = 1$ and $\alpha = 2$, respectively. Two parabolas are the normalized progress rates φ^* of the $(\mu/\mu_I, \lambda)$-ES and the $(\lambda)_{\text{opt}}$-ES. An intersection of φ^* curves and ψ lines represent the steady state point (stationary mutation strength s^*_{st} and stationary progress rate φ^*_{st}) at which the respective ES combinations works.

α is sufficiently large. Actually, it can even outperform the CSA version for a certain range of α values.

The comparison of the $(\lambda)_{\text{opt}}$-σSA-ES with the $(\mu/\mu_I, \lambda)$-ES shows that the $(\lambda)_{\text{opt}}$-σSA-ES outperform the simple multirecombination ES if the value of parameter α is sufficiently large. The reason is that $(\lambda)_{\text{opt}}$-σSA-ES uses the information about the ordering of all λ mutation vectors, while the $(\mu/\mu_I, \lambda)$-ES takes only the best μ offspring into account.

At the same time, the $(\mu/\mu_I, \lambda)$-ES performs better for small values of α. In order to find an explanation for this phenomenon, one can inspect the steady state condition (28) for the $(\mu/\mu_I, \lambda)$-ES and the $(\lambda)_{\text{opt}}$-σSA-ES graphically (Fig. 3). Intersections of the progress rate curves and the SAR lines in Fig. 3 correspond to the solutions of Eq. (28) for different values of α. For small α values ($\alpha = 1$ in Fig. 3), the steady state condition of the $(\mu/\mu_I, \lambda)$-ES holds true for the larger value of mutation strength compared to the $(\lambda)_{\text{opt}}$-σSA-ES. The steady state progress rate of the $(\mu/\mu_I, \lambda)$-ES is also larger for these values of the mutation strength. For larger α values ($\alpha = 2$ in Fig. 3), the steady state mutation strengths of the two ES versions are closer to each other, but the $(\lambda)_{\text{opt}}$-σSA-ES has a higher progress rate than the $(\mu/\mu_I, \lambda)$-ES.

The performances of the different ES versions are compared for different search space dimensionalities N in Fig. 4. The number of generations G was counted from first generation until the condition $f(\mathbf{y}) < 10^{-10}$ was satisfied. The initial conditions were $\mathbf{y}^{(0)} = \mathbf{1000}$ and $\sigma^{(0)} = 1$. For the $(\lambda)_{\text{opt}}$-σSA-ES with $\mu = 4$ and $\lambda = 10$, the weights have been chosen according to Eq. (22). The optimal α value $\alpha_{\text{opt}} = 4.6$, given by Eq. (37), was used. As to the $(\mu/\mu_I, \lambda)$-ES, α has been chosen according to Eq. (6), $\alpha = 0.7$.

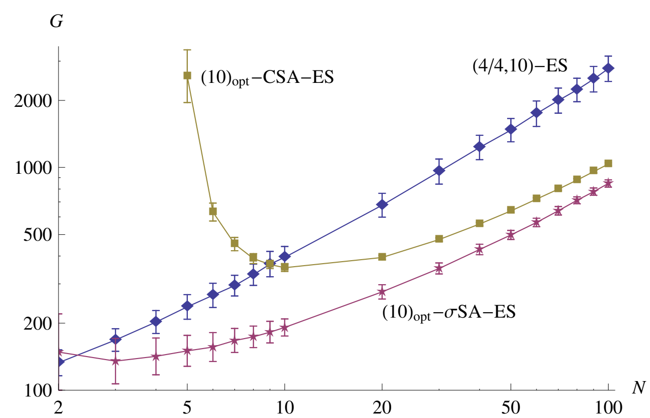

Figure 4: Number of generations G required to reach an objective function value of $f(\mathbf{y}) = 10^{-10}$ as a function of search space dimensionality N. Presented are the results of experiments averaged over 300 runs.

Comparing $(\lambda)_{\text{opt}}$-σSA-ES and $(\mu/\mu_I, \lambda)$-ES, it becomes immediately obvious that the new $(\lambda)_{\text{opt}}$-σSA-ES outperforms the old standard σSA-ES for almost all search space dimensionalities considered.

Somewhat surprisingly, the $(\lambda)_{\text{opt}}$-CSA-ES exhibits divergent behavior for $N < 5$, therefore no data are available for this strategy in Fig. 4 for low search space dimensionalities. A closer look at the ES dynamics reveals that the CSA is not able to perform mutation strength adaptation under these conditions. Instead, the mutation strength permanently decreases to zero. The deeper reason for the failure of the $(\lambda)_{\text{opt}}$-CSA-ES for small N is still an open question.

The $(\lambda)_{\text{opt}}$-σSA-ES is free from such divergence deficiency.

Furthermore, it demonstrates at the same time better performance for all N. As can be inferred from the plots in Fig. 4, this is not merely a theoretical result from the asymptotic theory, but a real observable behavior of the newly designed $(\lambda)_{\text{opt}}$-σSA-ES.

5. SUMMARY AND OUTLOOK

In this paper, the $(\lambda)_{\text{opt}}$-σSA-ES has been proposed as a new ES combining the strengths of Arnold's weighted recombination and of the mutative self-adaptive σ learning rule. The performance of this new strategy has been investigated on the sphere model. The theoretical analysis was carried out in the limit of infinite search space dimensionality. Real ES experiments were used in order to verify the predictive quality of the derived formulas in finite-dimensional search spaces. The experiments showed satisfactory agreement between the theoretical and experimental results which improves with increasing search space dimensionality. An experimental comparison of the performance showed that the $(\lambda)_{\text{opt}}$-σSA-ES outperforms the $(\lambda)_{\text{opt}}$-CSA-ES – which was regarded as the most efficient ES with isotropic mutations so far. Theoretical formulas derived in the limit of infinite search space dimensionality predicted the results of this comparison. This is also a strong argument for the usefulness of the dynamical systems approach in the field of evolutionary algorithm engineering. While we have shown that the newly designed $(\lambda)_{\text{opt}}$-σSA-ES works well on the sphere model, further investigations are necessary to evaluate its behavior on more complex test functions.

The choice of weights (22) discussed in this paper is not the only possible one. Alternatively, one can consider $\omega_{l,\lambda} = E_{l,\lambda}/\kappa$ where $\kappa > 1$. Using scaled weights, larger optimal $s^{*(g)}$ values can be obtained in the course of the mutation strength adaptation. While this does not increase the maximal possible quality gain in the non-noisy fitness case, it can be beneficial in noisy fitness environments and multimodal fitness landscapes as the strategy will work with larger mutations [5]. The question how the choice of scaled weights influences the performance in non-noisy and noisy fitness environments remains to be investigated in future research. To this end, empirical and theoretical investigations are to be conducted considering more complex test functions such as PDQFs [10, 2] and general quadratic models including noise.

6. ACKNOWLEDGMENTS

This work was supported by the Austrian Science Fund (FWF) under grant P19069-N18.

7. REFERENCES

[1] D. V. Arnold. Weighted multirecombination evolution strategies. Theoretical computer science, 361(1):18–37, 2006.

[2] D. V. Arnold. On the use of evolution strategies for optimising certain positive definite quadratic forms. In GECCO '07: Proceedings of the 9th annual conference on Genetic and evolutionary computation, pages 634–641, New York, NY, USA, 2007. ACM.

[3] D. V. Arnold and H.-G. Beyer. Local Performace of the $(\mu/\mu_I, \lambda)$-ES in a Noisy Environment. In W. Martin and W. Spears, editors, Foundations of Genetic Algorithms, 6, pages 127–141, San Francisco, CA, 2001. Morgan Kaufmann.

[4] T. Bartz-Beielstein. Experimental Research in Evolutionary Computation - The New Experimentalism. Natural Computing Series. Springer, Berlin, 2006.

[5] H.-G. Beyer. The Theory of Evolution Strategies. Natural Computing Series. Springer, Heidelberg, 2001.

[6] H.-G. Beyer. Evolution Strategies. Scholarpedia, page 15321, 2007.

[7] H.-G. Beyer and D. Arnold. Qualms Regarding the Optimality of Cumulative Path Length Control in CSA/CMA-Evolution Strategies. Evolutionary Computation, 11(1):19–28, 2003.

[8] L. Grünz and H.-G. Beyer. Some Observations on the Interaction of Recombination and Self-Adaptation in Evolution Strategies. In P. Angeline, editor, Proceedings of the CEC'99 Conference, pages 639–645, Piscataway, NJ, 1999. IEEE.

[9] N. Hansen and A. Ostermeier. Completely Derandomized Self-Adaptation in Evolution Strategies. Evolutionary Computation, 9(2):159–195, 2001.

[10] J. Jägersküpper. How the (1+1)-ES Using Isotropic Mutations Minimizes Positive Definite Quadratic Forms. Theoretical Computer Science, 361(1):38–56, 2006.

[11] S. Meyer-Nieberg. Self-Adaptation in Evolution Strategies. PhD thesis, University of Dortmund, CS Department, Dortmund, Germany, 2007.

[12] S. Meyer-Nieberg and H.-G. Beyer. On the Analysis of Self-Adaptive Recombination Strategies: First Results. In Proceedings of the CEC'05 Conference, pages 2341–2348, Piscataway, NJ, 2005. IEEE.

[13] A. Ostermeier, A. Gawelczyk, and N. Hansen. A Derandomized Approach to Self-Adaptation of Evolution Strategies. Evolutionary Computation, 2(4):369–380, 1995.

[14] I. Rechenberg. Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution. Frommann-Holzboog Verlag, Stuttgart, 1973.

[15] H.-P. Schwefel. Numerical Optimization of Computer Models. Wiley, Chichester, 1981.